

What is claimed is:

1. A method of operating a file server for creating and maintaining a read-write snapshot copy of a production file system, the method comprising:

5 a) creating a read-only snapshot copy of the production file system; and

b) creating a read-write snapshot copy of the production file system based on the read-only snapshot copy of the production file system, and maintaining the read-write snapshot copy of the production file system by maintaining a set of save volume blocks of new data of the read-write snapshot copy of the production file system, and

10 (i) writing to a first specified block of the read-write snapshot copy of the production file system by writing to a respective save volume block allocated to the first specified block of the read-write snapshot copy of the production file system,

(ii) reading from a second specified block of the read-write snapshot copy of the production file system by checking whether there is a respective save volume block allocated to the second specified block of the read-write snapshot copy of the production file system, and upon finding that there is a respective save volume block allocated to the second specified block of the read-write snapshot copy of the production file system, reading from the respective save volume block found to be allocated to the second specified block of the read-write snapshot copy of the production file system, and

(iii) reading from a third specified block of the read-write snapshot copy of the production file system by checking whether there is a respective save

volume block allocated to the third specified block of the read-write snapshot copy of the production file system, and upon finding that there is not a respective save volume block allocated to the third specified block of the read-write snapshot copy of the production file system, reading the third specified block of the read-write snapshot copy of the production file system by reading from the read-only snapshot copy of the production file system.

2. The method as claimed in claim 1, wherein the respective save volume block allocated to the first specified block of the read-write snapshot copy of the production file system is allocated in response to the file server receiving a request from a client for writing to the first specified block of the read-write snapshot copy of the production file system.

3. The method as claimed in claim 1, which includes the file server maintaining a bit map for indicating whether or not respective blocks of the read-write snapshot copy of the production volume have new data residing in save volume blocks, and maintaining a block map for indicating respective save volume addresses of blocks of new data of the read-write snapshot copy of the production file system.

4. The method as claimed in claim 1, which includes creating and maintaining more than one read-write snapshot copy of the production file system based on the read-only snapshot copy of the production file system.

5. The method as claimed in claim 1, wherein the read-write snapshot copy of the production file system has a file system identifier, and which includes refreshing the read-write snapshot copy of the production file system by freezing access to the read-write snapshot copy of the production file system, de-allocating the set of save volume
5 blocks for storing new data of the read-write snapshot copy of the production file system, creating a new read-only snapshot copy of the production file system, creating a new read-write snapshot copy of the production file system based on the new read-only snapshot copy of the production file system, and assigning the file system identifier to the new read-write snapshot copy of the production file system.

10

6. The method as claimed in claim 1, which includes restoring the production file system from the read-write snapshot copy of the production file system by freezing the production file system, marking the production file system as being under restoration, creating a new read-only snapshot copy of the production file system including a new set
15 of save volume blocks and a new bitmap for identifying blocks written to the production file system since the time of restoring the production file system from the read-write snapshot copy of the production file system, launching a background process of copying, either to the production file system or to the new set of save volume blocks as indicated by the new bitmap, save volume blocks of data of the read-write snapshot copy of the
20 production file system, thawing the production file system to permit read-write access of the production file system under a foreground routine, and when the background process is done, returning the production file system to normal read-write access.

7. A method of operating a file server for providing access to a production file system and a read-write snapshot copy of the production file system, the method comprising:

the file server maintaining a clone volume of blocks in the production file system,
5 a set of save volume blocks of old data from the production file system, and a set of save volume blocks of new data of the read-write snapshot copy file system,

the file server writing new data to a first specified block in the production file system by copying old data from the first specified block in the clone volume to a save volume block and then writing to the first specified block in the clone volume,

10 the file server reading a second specified block of the production file system by reading the second specified block from the clone volume,

the file server writing new data to a third specified block in the read-write snapshot copy of the production file system by writing to a save volume block, and

the file server reading a fourth specified block from the read-write snapshot copy
15 of the production file system by checking whether the fourth specified block is found in the set of save volume blocks of new data of the read-write snapshot copy of the production file system, and upon finding that the fourth specified block is found in the set of save volume blocks of new data of the read-write snapshot copy of the production file system, reading the fourth specified block from the set of save volume blocks of new data
20 of the read-write snapshot copy of the production file system.

8. The method as claimed in claim 7, which includes the file server reading a fifth specified block from the read-write snapshot copy of the production file system by

checking whether the fifth specified block is found in the set of save volume blocks of new data of the read-write snapshot copy of the production file system, and upon finding that the fifth specified block is not found in the set of save volume blocks of new data of the read-write snapshot copy of the production file system, checking whether the fifth
5 specified block is found in the set of save volume blocks of old data from the production file system, and upon finding that the fifth specified block is found in the set of save volume blocks of old data from the production file system, reading the fifth specified block from the set of save volume blocks of old data from the production file system.

10 9. The method as claimed in claim 7, which includes the file server reading a sixth specified block from the read-write snapshot copy of the production file system by checking whether the sixth specified block is found in the set of save volume blocks of new data of the read-write snapshot copy of the production file system, and upon finding that the sixth specified block is not found in the set of save volume blocks of new data of
15 the read-write snapshot copy of the production file system, checking whether the sixth specified block is found in the set of save volume blocks of old data from the production file system, and upon finding that the sixth specified block is not found in the set of save volume blocks of old data from the production file system, reading the sixth specified block from the clone volume.

20

10. The method as claimed in claim 7, which includes the file server storing, in a shared save volume, the set of save volume blocks of old data from the production file

system and the set of save volume blocks of new data of the read-write snapshot copy of the production file system.

11. The method as claimed in claim 7, which includes the file server maintaining a first bit map for indicating whether or not respective blocks of the production volume have old data residing in the set of save volume blocks for storing old data from the production file system, maintaining a first block map for indicating respective save volume addresses of blocks of old data from the production file system, maintaining a second bit map for indicating whether or not respective blocks of the read-write snapshot copy of the production volume have new data residing in the set of save volume blocks of new data of the read-write snapshot copy of the production volume, and maintaining a second block map for indicating respective save volume addresses of blocks of new data of the read-write snapshot copy of the production file system.

12. The method as claimed in claim 7, wherein the read-write snapshot copy of the production file system has a file system identifier, and wherein the file server refreshes the read-write snapshot copy of the production file system by freezing access to the read-write snapshot copy of the production file system, de-allocating the set of save volume blocks for storing new data of the read-write snapshot copy of the production file system, creating a new read-only snapshot copy of the production file system, creating a new read-write snapshot copy of the production file system based on the new read-only snapshot copy of the production file system, and assigning the file system identifier to the new read-write snapshot copy of the production file system.

13. The method as claimed in claim 7, which includes the file server restoring the production file system from the read-write snapshot copy of the production file system by freezing the production file system, marking the production file system as being under
5 restoration, creating a new read-only snapshot copy of the production file system including a new save volume and a new bitmap for identifying blocks written to the production file system since the time of restoring the production file system from the read-write snapshot copy of the production file system, launching a background process of copying, to either the clone volume or to the new save volume as indicated by
10 respective bits in the new bitmap, save volume blocks of data of the read-write snapshot copy of the production file system that are not in the clone volume or in the new save volume, thawing the production file system to permit read-write access of the production file system under a foreground routine, and when the background process is done, returning the production file system to normal read-write access.

15

14. The method as claimed in claim 13, which includes the file server performing the background process of copying save volume blocks by first copying the set of save volume blocks of old data from the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, and then
20 copying the set of save volume blocks of new data of the read-write snapshot copy of the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap.

15. The method as claimed in claim 13, which includes file server performing the background process of copying save volume blocks by first copying the set of save volume blocks of new data of the read-write snapshot copy of the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, and then, for production file system blocks not already copied from the set of save volume blocks of new data of the read-write snapshot copy of the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, copying save volume blocks of old data from the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap.

16. A file server comprising:
means for storing a production file system;
means for creating and maintaining a series of read-only snapshot copies of the production file system; and
means for creating and maintaining read-write snapshot copies of the production file system, each of the read-write snapshot copies of the production file system being based on a respective one of the read-only snapshot copies of the production file system.

17. The file server as claimed in claim 16, which includes means for refreshing a selected read-write snapshot copy of the production file system with a specified read-only snapshot copy of the production file system.

18. The file server as claimed in claim 16, which includes means for refreshing a selected read-write snapshot copy of the production file system by creating a new read-only snapshot copy of the production file system and refreshing the selected read-write snapshot copy of the production file system with the new read-only snapshot copy of the
5 production file system.

19. The file server as claimed in claim 16, which includes means for restoring the production file system from a specified read-write snapshot copy of the production file system.

10

20. A file server for creating and maintaining a read-write snapshot copy of a production file system, the file server being programmed for creating a read-only snapshot copy of the production file system, creating a read-write snapshot copy of the production file system based on the read-only snapshot copy of the production file
15 system, and maintaining the read-write snapshot copy of the production file system by maintaining a set of save volume blocks of new data of the read-write snapshot copy of the production file system,

wherein the file server is programmed for writing to a first specified block of the read-write snapshot copy of the production file system by writing to a respective save
20 volume block allocated to the first specified block of the read-write snapshot copy of the production file system, and

wherein the file server is programmed for reading from a second specified block of the read-write snapshot copy of the production file system by checking whether there

is a respective save volume block allocated to the second specified block of the read-write snapshot copy of the production file system, and upon finding that there is a respective save volume block allocated to the second specified block of the read-write snapshot copy of the production file system, reading from the respective save volume
5 block found to be allocated to the second specified block of the read-write snapshot copy of the production file system, and

wherein the file server is programmed for reading from a third specified block of the read-write snapshot copy of the production file system by checking whether there is a respective save volume block allocated to the third specified block of the read-write
10 snapshot copy of the production file system, and upon finding that there is not a respective save volume block allocated to the third specified block of the read-write snapshot copy of the production file system, reading the third specified block of the read-write snapshot copy of the production file system by reading from the read-only snapshot copy of the production file system.

15

21. The file server as claimed in claim 20, which is programmed for allocating the respective save volume block to the first specified block of the read-write snapshot copy of the production file system in response to the file server receiving a request from a client for writing to the first specified block of the read-write snapshot copy of the
20 production file system.

22. The file server as claimed in claim 20, wherein the file server is programmed for maintaining a bit map for indicating whether or not respective blocks of the read-write

snapshot copy of the production volume have new data residing save volume blocks, and for maintaining a block map for indicating respective save volume addresses of blocks of new data of the read-write snapshot copy of the production file system.

- 5 23. The file server as claimed in claim 20, which is programmed for creating and maintaining more than one read-write snapshot copy of the production file system based on the read-only snapshot copy of the production file system.

- 10 24. The file server as claimed in claim 20, wherein the read-write snapshot copy of the production file system has a file system identifier, and wherein the file server is programmed for refreshing the read-write snapshot copy of the production file system by freezing access to the read-write snapshot copy of the production file system, de-allocating the set of save volume blocks for storing new data of the read-write snapshot copy of the production file system, creating a new read-only snapshot copy of the
15 production file system, creating a new read-write snapshot copy of the production file system based on the new read-only snapshot copy of the production file system, and assigning the file system identifier to the new read-write snapshot copy of the production file system.

- 20 25. The file server as claimed in claim 20, which is programmed for restoring the production file system from the read-write snapshot copy of the production file system by freezing the production file system, marking the production file system as being under restoration, creating a new read-only snapshot copy of the production file system

including a new set of save volume blocks and a new bitmap for identifying blocks written to the production file system since the time of restoring the production file system from the read-write snapshot copy of the production file system, launching a background process of copying, either to the production file system or to the new set of save volume blocks as indicated by the new bitmap, save volume blocks of data of the read-write snapshot copy of the production file system, thawing the production file system to permit read-write access of the production file system under a foreground routine, and when the background process is done, returning the production file system to normal read-write access.

10

26. A file server for providing access to a production file system and a read-write snapshot copy of the production file system,

the file server being programmed to maintain a clone volume of blocks in the production file system, a set of save volume blocks of old data from the production file system, and a set of save volume blocks of new data of the read-write snapshot copy file system,

15

the file server being programmed to write new data to a first specified block in the production file system by copying old data from the first specified block in the clone volume to a save volume block and then writing to the first specified block in the clone volume,

20

the file server being programmed to read a second specified block of the production file system by reading the second specified block from the clone volume,

the file server being programmed to write new data to a third specified block in the read-write snapshot copy of the production file system by writing to a save volume block, and

the file server being programmed to read a fourth specified block from the read-write snapshot copy of the production file system by checking whether the fourth specified block is found in the set of save volume blocks of new data of the read-write snapshot copy of the production file system, and if so, reading the fourth specified block from the set of save volume blocks of new data of the read-write snapshot copy of the production file system, and if not, checking whether the fourth specified block is found in the set of save volume blocks of old data from the production file system, and if so, reading the fourth specified block from the set of save volume blocks of old data from the production file system, and if not, reading the fourth specified block from the clone volume.

27. The file server as claimed in claim 26, wherein the set of save volume blocks of old data from the production file system and the set of save volume blocks of new data of the read-write snapshot copy of the production file system are contained in a shared save volume.

28. The file server as claimed in claim 26, which is programmed for maintaining a first bit map for indicating whether or not respective blocks of the production volume have old data residing in the set of save volume blocks for storing old data from the production file system, maintaining a first block map for indicating respective save

volume addresses of blocks of old data from the production file system, maintaining a second bit map for indicating whether or not respective blocks of the read-write snapshot copy of the production volume have new data residing in the set of save volume blocks of new data of the read-write snapshot copy of the production volume, and maintaining a
5 second block map for indicating respective save volume addresses of blocks of new data of the read-write snapshot copy of the production file system.

29. The file server as claimed in claim 26, wherein the read-write snapshot copy of the production file system has a file system identifier, and wherein the file server is
10 programmed for refreshing the read-write snapshot copy of the production file system by freezing access to the read-write snapshot copy of the production file system, de-allocating the set of save volume blocks for storing new data of the read-write snapshot copy of the production file system, creating a new read-only snapshot copy of the production file system, creating a new read-write snapshot copy of the production file
15 system based on the new read-only snapshot copy of the production file system, and assigning the file system identifier to the new read-write snapshot copy of the production file system.

30. The file server as claimed in claim 26, wherein the file server is programmed for
20 restoring the production file system from the read-write snapshot copy of the production file system by freezing the production file system, marking the production file system as being under restoration, creating a new read-only snapshot copy of the production file system including a new save volume and a new bitmap for identifying blocks written to

the production file system since the time of restoring the production file system from the read-write snapshot copy of the production file system, launching a background process of copying, to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, save volume blocks of data of the read-write snapshot
5 copy of the production file system that are not in the clone volume or in the new save volume, thawing the production file system to permit read-write access of the production file system under a foreground routine, and when the background process is done, returning the production file system to normal read-write access.

10 31. The file server as claimed in claim 30, wherein the file server is programmed to perform the background process of copying save volume blocks by first copying the set of save volume blocks of old data from the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, and then copying the set of save volume blocks of new data of the read-write snapshot copy
15 of the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap.

32. The file server as claimed in claim 30, wherein the file server is programmed to perform the background process of copying save volume blocks by first copying the set
20 of save volume blocks of new data of the read-write snapshot copy of the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, and then, for production file system blocks not already copied from the set of save volume blocks of new data of the read-write snapshot copy of the

production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap, copying save volume blocks of old data from the production file system to either the clone volume or to the new save volume as indicated by respective bits in the new bitmap.